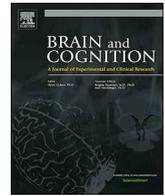




ELSEVIER

Contents lists available at ScienceDirect

Brain and Cognition

journal homepage: www.elsevier.com/locate/b&c

Embodying addiction: A predictive processing account

Mark Miller^{a,*}, Julian Kiverstein^{b,e}, Erik Rietveld^{b,c,d,e}

^a Department of Informatics, University of Sussex, Sussex House, Falmer, Brighton BN1 9RH, United Kingdom

^b Amsterdam University Medical Center, Department of Psychiatry, University of Amsterdam, Amsterdam, the Netherlands

^c Department of Philosophy, Institute for Logic, Language and Computation, University of Amsterdam, the Netherlands

^d Department of Philosophy, University of Twente, Enschede, the Netherlands

^e Amsterdam Brain and Cognition Centre, University of Amsterdam, Amsterdam, the Netherlands

ABSTRACT

In this paper we show how addiction can be thought of as the outcome of learning. We look to the increasingly influential predictive processing theory for an account of how learning can go wrong in addiction. Perhaps counter intuitively, it is a consequence of this predictive processing perspective on addiction that while the brain plays a deep and important role in leading a person into addiction, it cannot be the whole story. We'll argue that predictive processing implies a view of addiction not as a brain disease, but rather as a breakdown in the dynamics of the wider agent-environment system. The environment becomes meaningfully organised around the agent's drug-seeking and using behaviours. Our account of addiction offers a new perspective on what is harmful about addiction. Philosophers often characterise addiction as a mental illness because addicts irrationally shift in their judgement of how they should act based on cues that predict drug use. We argue that predictive processing leads to a different view of what can go wrong in addiction. We suggest that addiction can prove harmful to the person because as their addiction progressively takes hold, the addict comes to embody a predictive model of the environment that fails to adequately attune them to a volatile, dynamic environment. The use of an addictive substance produces illusory feedback of being well-attuned to the environment when the reality is the opposite. This can be comforting for a person inhabiting a hostile niche, but it can also prove to be harmful to the person as they become skilled at living the life of an addict, to the neglect of all other alternatives. The harm in addiction we'll argue is not to be found in the brains of addicts, but in their way of life.

1. Introduction

Addiction has a devastating effect upon those whose life it afflicts. Addicts find their life increasingly dominated by their addictive behaviours. The other pursuits they care about begin to be crowded out as they devote increasing amounts of time and energy to servicing their addiction. The undesirable outcomes of their addictive behaviours are increasingly ignored by them, yet at the same time addicts feel compelled to continue acting on their addictions often long after the addictive behaviour has ceased to bring any pleasure. Addiction can reach a point in a person's life where it seems all that matters to them is doing what their addiction demands of them, yet at the same time this is also often something they do not want. The director of the National Institute on Drug Abuse, Nora Volkow, has observed "I've never come across a single person that was addicted that wanted to be addicted" (Gugliotta, 2003).

Volkow's claim she has never encountered a willing addict is perhaps something of an exaggeration. Flanagan (2017) has shown through a range of convincing examples how some addicts may reasonably prefer, all things considered, to remain addicts. These individuals may legitimately be said to want their addictive lifestyles to continue, more or less unchanged. This needn't be because they

compare the costs of stopping using a substance with continuing, and conclude the best course of action is to preserve the status quo. It maybe because using a drug helps them numb physical or emotional pain, or to have the novel and meaningful experiences they can only achieve through the use of a substance, or because using a substance gives them a sense of group belonging (Flanagan, 2017: p.68). More generally, an individual's cultural and economic circumstances may allow them to avoid incurring the personal, social and economic costs addiction typically incurs. They may not experience any of the shame, depression and loss addicts often suffer.

What does seem right in Volkow's observation however is that many addicts see the harm that a continued use of a substance will do to themselves and those around them, and desire more than anything to change. Yet they also feel a strong compulsion to continue using a drug or to drink, and find themselves again doing what their addiction demands of them. They are unwilling addicts; their lifestyle is one that causes them to suffer yet they feel powerless to change the life they lead. This raises the central questions that will occupy us in this paper: When does the behaviour of an addict cross the line from contributing to the person's well-being to being harmful to them? When should the habits an addict develops be thought of as bad habits?

In what follows we will argue *addiction should be thought of as a self-*

* Corresponding author.

E-mail addresses: m.d.miller@sussex.ac.uk (M. Miller), j.d.kiverstein@amsterdamumc.nl (J. Kiverstein), d.w.rietveld@amsterdamumc.nl (E. Rietveld).

<https://doi.org/10.1016/j.bandc.2019.105495>

Received 16 September 2019; Received in revised form 8 November 2019; Accepted 13 November 2019

Available online 23 December 2019

0278-2626/ © 2019 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license

(<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

organising process of a whole agent-environment system. Addiction becomes harmful to a person when this self-organising process spirals out of control due to feedback loops that entrain the behaviour of the agent, locking them into destructive cycles of behaviour. It is in the dynamic interaction between the agent and its environment that addiction is born and endures.

In Section 2 we review the effects of substances of addiction on dopaminergic neurons in the midbrain.¹ Repeated use of drugs has been shown to induce functional and structural changes in the nervous systems of addicts by acting either directly or indirectly on dopaminergic neurons. This finding has led the medical community, and many scientists, to embrace a “disease model” of addiction. According to the disease model, addicts find it difficult to change their behaviour even when this is what they desire because of dysfunctional neural systems that mediate reward learning. Reward learning has wide-reaching functional influences, mediating everything from perception and memory, emotion and attention, to decision-making and cognitive control. It is the effect of drugs on reward learning that the disease model takes to explain compulsive drug seeking and use on the one hand. The long-term functional and structural changes prolonged substance abuse induces have in turn been taken to explain why addicts find it so hard to change their behaviour, and frequently relapse.

A new picture of reward learning is however beginning to take shape in current cognitive neuroscience according to which reward is the consequence of action, not its cause. Rewarding outcomes are predicted, and the agent selects actions that fulfill its predictions (den Ouden, Daunizeau, Roiser, Friston, & Stephan, 2010; Friston, Daunizeau, & Kiebel, 2009; Clark, 2015a: ch. 4; FitzGerald, Dolan, & Friston, 2014; Friston et al., 2012). In Section 3 we provide an overview of this new predictive processing perspective on reward learning. In the predictive processing, or active inference, account of action and perception, dopaminergic discharges are best thought of as weighing the agent’s confidence in relevant affordances given their skills and abilities (Bruineberg, Kiverstein, & Rietveld, 2018; Friston et al., 2012; Kiverstein, Miller, & Rietveld, 2019; Kiverstein, Rietveld, Slagter, & Denys, 2019; Linson, Clark, Ramamoorthy, & Friston, 2018). In brief, active inference is the process of selecting the affordances that stand out as relevant or inviting because they are expected to minimise long-term prediction error - the mismatch between expected and current sensory inputs on average, and in the long-run. We characterise active inference as the process of selecting those affordances that are relevant to the agent because they are likely to lead to expected sensory outcomes (Kiverstein, Miller, et al., 2019; Kiverstein, Rietveld, et al., 2019). Some affordances are more likely than others to lead to expected, unsurprising outcomes. The agent should therefore be selectively open to those affordances that have the highest probability of leading from their current situation to expected outcomes. This implies a probability distribution over relevant affordances that can have a greater or lesser precision (i.e., salience or reliability). Substantial evidence suggests that dopamine scores *precision* - the agent’s confidence about what affordances will take them their current sensory states to their expected future sensory outcomes given what they are capable of doing and the current context.²

The disease model claims that what is pathological in addiction is the way in which dopaminergic processes in the brain are hijacked by substances of addiction. The behaviour of the addict comes to be passively and automatically driven by sensory cues that predict drug use, and is progressively less and less under the control of their conscious evaluation. In Section 4 we argue this is the wrong conclusion to draw

by showing how the predictive processing theory supports an ecological-enactive account of habits. We go on in Section 5 to put this account of habits to work to explain what can go badly wrong in substance addiction. What can prove harmful about addiction is, we suggest, the build-up of error over time, and a failure to adequately assign precision to relevant affordances in the addict’s dealings with a dynamic, and volatile world. In Section 6 we draw on our earlier work on the role of what we call “error dynamics” in the context-sensitive weighing of precision (Kiverstein, Miller, et al., 2019). Error dynamics refer to the *rate of change* in prediction error over time. We show how error dynamics are sensed by the agent in the form of positively or negatively valenced bodily feelings. Bodily feelings track whether the agent is doing better or worse than expected at minimising error in their engagement with the environment. We hypothesise that dopaminergic systems are among the systems in the brain that track error dynamics. In Section 7 we argue that when substances of addiction act directly on these systems, they provide false feedback that the agent is doing better than expected at minimising error, when the reality is often the opposite. Section 8 concludes our argument by returning to the disease model. We show how addiction is not a disease of the brain, but needs to be understood in the context of the wider dynamic of the agent’s coupling to its environment. The brain is of course a necessary part of this story, but it isn’t sufficient for understanding what can go wrong in addiction. Addiction can prove to be harmful to the person as they become skilled at living the life of an addict, to the neglect of all other alternatives. The harm in addiction we’ll argue is thus not to be found in the brains of addicts, but in their way of life.

2. Mutiny in the midbrain: Is addiction a brain disease?

The disease model of addiction claims that addiction is a “chronic, relapsing brain disease that is characterised by compulsive drug seeking and use, despite harmful consequences” (National Institute on Drug Abuse, 2009; World Health Organisation, 2004). Addiction is conceived of as pathological because substance use leads the reward learning system in the brain to malfunction, leading the person to compulsively seek out and use a substance despite the negative consequences of doing so. The key idea behind this model of substance addiction is that substance use leads to a “hijacking” of reward learning systems in the brain (Ahmed, 2004; Schultz, 2016; Keramati & Gutkin, 2013; c.f. Elster, 1999; Gardner & David, 1999; Redish, 2004; Everitt et al., 2008).³ Reward-based learning is standardly understood as the process by which the organism maximises expected utility while minimising costs and avoiding punishment (Arpaly & Schroeder, 2013; Delgado, Miller, Inati, & Phelps, 2005; Sutton & Barto, 1998). Reward learning steers agents through the world in ways that increase the probability of finding what is valuable and avoiding what is aversive or punishing to the agent. Agents learn about values (e.g. expected rewards) by means of “reward prediction error” (RPE) signals. These signals are modelled as computing the difference between received and predicted rewards.⁴

³ There is disagreement in the literature about whether addiction is a pathology of reward learning or of motivation. Berridge and colleagues have shown how rats deprived of dopamine through a lesioning of their mesolimbic system continued to learn the reward value of a stimulus but are no longer motivated to act on this learning. They are no longer prepared to work (say pressing a bar) to attain a food reward (see e.g. Berridge, 2007; Holton & Berridge, 2013). Berridge and colleagues propose an incentive salience model of the dopamine system. They argue that the dopamine system directly causes desires that drive action. Addiction is thus a pathology of wanting or desire, not of learning. We argue for a new version of the learning accounts below. We will show how Berridge’s notion of incentive salience maps onto what is referred to as “precision” in the predictive processing theory. Thus, our predictive processing theory of additions hold the promise of reconciling the learning and motivation theories of addiction.

⁴ Dopamine neurons in midbrain areas have been found to fire at rates that

¹ We will focus on substance addiction in what follows. While we believe some features of our account may generalise to other forms of addiction, this is not something we explore in this paper.

² Our thanks to an anonymous reviewer for help with wording here in our characterisation of active inference.

The job of signalling unexpected reward is hypothesised to be performed in the brain by the phasic bursts of dopamine (Montague, Dayan, & Sejnowski, 1996; Schultz, Dayan, & Montague, 1997). Mesolimbic neurons in the midbrain fire when an unexpected reward is delivered. Mammals and invertebrates alike respond with reward learning signals when presented with unexpected opportunities for reward such as food, water and mates (Schultz et al., 1997; Sutton & Barto, 1998). The midbrain regions in the human brain respond in a similar way to such things as money, success, favorite songs, and the flourishing of loved ones (Arpaly & Schroeder, 2013).

There is now a substantial body of research that supports the claim that substances of addiction act on reward learning circuitry in ways that lead to changes in the function and structure of the addict's brain. Drug use enhances the formation of new synapses, strengthening connections between the striatum; amygdala; and hippocampus, while at the same time reducing synaptic density in prefrontal cortex (Goldstein & Volkow, 2011; Volkow, Koob, & McLellan, 2016). These changes in the brain's wiring are correlated with reduced capacity to engage in cognitive control, compulsivity in drug seeking, and blunting of reward response more generally (Everitt & Robbins, 2013; Volkow & Fowler, 2000). However, it is worth remarking that every experience that is repeated sufficiently many times will induce comparable changes in brain wiring to those seen in the users of addictive substances (Levy, 2013; Lewis, 2015). The development of any habit, good or bad, will have many of the same signature kinds of changes in brain wiring that are seen in addiction. The disease model looks to such neurochemical alterations to answer the question of what can go wrong in addiction. However, given that many of the same changes in neural function and structure are seen across the board in the development of habits, the disease model fails to identify the neurocognitive mechanisms that lead to harmful behaviour in addicts. Marc Lewis makes this point well when he observes:

“Addiction” doesn't fit a unique physiological stamp. It simply describes the repeated pursuit of highly attractive goals, and the brain changes that condense this cycle of thought and behaviour into a well-learned habit.” (Lewis, 2017: p.12)

The disease model of addiction tends to conceive of the learning that leads to addiction in terms of Pavlovian conditioning. The strong compulsion to use the substance is thought to be due to the dorsal striatum no longer being under the control of cortical areas (such as the dorsolateral prefrontal cortex) believed to ordinarily regulate and contextualize habitual responses (Everitt & Robbins, 2013). Environmental cues come to passively drive behaviour, eliciting powerful urges to consume the substance independent of what the person wants. In the next section we will sketch a different picture that is beginning to emerge of the role of dopamine in learning. Instead of signalling reward prediction error, dopamine is modelled as signalling confidence about how to act in the world. We do not dispute the evidence that dopamine has a central role to play in a person's developing addictive patterns of behaviour. But we suggest the role of dopamine may be somewhat different from how it is standardly understood in the disease model. Instead of the person being driven to act passively on the basis of cues that predict reward, dopamine contributes to attuning the person to the affordances that are relevant to them in the environment. We will argue this shift in perspective in computational neuroscience should lead to a different view of what is harmful in addiction. Dopamine should be understood as tuning the agent to possibilities for action that allow them to flourish in the long run. But drugs of addiction, while masquerading as increasing attunement to what is important to the person,

(footnote continued)

directly correlate with RPE signals and perceived reward values providing support for the hypothesis that dopamine functions as a learning signal in the brain (Montague et al., 1996; Schultz et al., 1997).

can in fact lead to greater disattunement over time. Thus, instead of thinking of addiction as a disorder of the brain, we should conceive of the harm addiction does as resulting from increasing disorder in the agent-environment system the person forms with their ecological niche.

3. The predictive processing theory of reward learning

In recent years an account of reward learning has begun to take shape that takes predictions of reward to be part and parcel of the prediction of the sensory consequences of our acting in the world (Clark, 2015b; FitzGerald et al., 2014; Friston et al., 2012; Friston et al., 2009). It seems intuitive to think of maximising utility as a complex cause of an agent's behaviour. Agents should be motivated by acquiring rewards and avoiding punishments. They should, when all goes well, thereby learn over time to frequent rewarding spaces more often than not (and to avoid punishing spaces whenever possible). The predictive processing theory (PP) we review in this section turns this intuition about expectation and reward on its head (FitzGerald et al., 2014; Friston et al., 2012). PP starts from expectations, and not from rewards, as the causes of behaviour. Agents act to minimise surprise about their own future sensory states. It is future proprioceptive, interoceptive and exteroceptive states associated with a course of action that are predicted. “Surprise” in the relevant technical sense (also referred to as “surprisal”) thus relates to expected future sensory states. Surprise corresponds to self-information; namely, the implausibility of some action outcome on average and over time. Desired future sensory states are more likely, and thus less surprising, than undesired states. Thus in relation to future sensory states (action outcomes), we have expected surprise (i.e., expected self-information)⁵, which is uncertainty (i.e., expected self-information). In short, agents act to minimise uncertainty in relation to their engagement with a field of multiple relevant affordances.⁶ If I want a coffee, the unsurprising outcome would be for me to find myself in the near future in the cafe were I typically buy my favourite coffee. To minimise surprise (i.e. long-term prediction error) then the agent must select the actions that are most likely to lead them from their current sensory states to those they expect to occupy in the future. It is therefore unsurprising or expected outcomes that are rewarding. Reward is the consequence of behaviour, not its cause - the agent is rewarded when the future outcomes of its actions are expected or unsurprising.

In Section 2 we've seen how drugs can act directly or indirectly on midbrain areas, increasing the transmission of dopamine from the midbrain to the forebrain structures. The PP theory understands these effects on the dopamine system as impacting on the optimisation of what are called “precision expectations”. Before we can see how this might work, we must briefly explain the notion of precision.

To succeed in minimising future prediction errors through action, an agent must have some means of determining its own uncertainty about the effects of its actions. An agent that minimises surprise should act in ways that tend over time to minimise the divergence or the difference between attainable and expected outcomes (Friston et al., 2014). We will use the term “precision” to refer to the confidence associated with relevant affordances. Precision marks the agent's confidence that given their skills and abilities, affordances can take them from their current sensory state to expected outcomes that are

⁵ Many thanks to an anonymous reviewer for pointing out the connection of surprise and self-information. “Self-information” is a concept from information theory that refers to the amount of information that is gained from the sampling of a random sensory signal.

⁶ We stress that the characterisation of active inference and predictive processing in terms of engagement with multiple relevant affordances is our own interpretation of active inference, although we believe many of the papers by the original architects of the active inference theory can be interpreted in these terms without distortion (see e.g. Friston, 2011; Friston et al., 2012; Linson et al., 2018).

unsurprising, and thus rewarding.⁷ Low-precision on a relevant affordance means it is likely that future sensory states will fail to match with those that are expected, and therefore the affordance should not invite the agent to act. A surprise minimizing agent should therefore allow such affordances to have only a minimal influence in the regulation of its behaviour. High precision relevant affordances by contrast have a high probability of leading to a match between expected or desired outcomes, and the future sensory states the agent attains as the effects of its actions. The predictions of future sensory states whose precision is weighted we will call “action policies”. An action policy refers to a sequence of actions - a path that takes the agent from its current sensory states to those it expects to occupy. Our policy of frequenting our favourite cafe to buy coffee is an example of such a high-precision policy. We have high confidence that acting on such a policy is likely to minimise the difference between the sensory states we predict ourselves occupying when we visit this cafe, and the sensory states we expect or desire - those associated with drinking an excellent mug of coffee.

In PP the midbrain dopamine system is assigned the function of weighing the confidence about what to do next (Friston, 2012). The firing of dopaminergic neurons in this account doesn't report reward prediction error, but rather a “salient and unexpected event, under varying degrees of ambiguity or uncertainty” (Friston, 2012: p.277). Precision isn't something that is known in advance, but has to be learned. The dopamine system is hypothesised to be a part of the machinery for optimising precision expectations through learning.⁸ The precision of a relevant affordance depends on how attainable an expected future sensory state is from current sensory states. Dopamine signals the likelihood that current sensory information anticipates a predictable sequence of actions (Friston et al., 2014; Linson et al., 2018).

Substance abuse leads to the learning of sub-optimal precision expectations (Schwartenbeck et al., 2015). The agent comes to place too much confidence in its top-down predictions of future sensory states for policies of drug-seeking and drug-using behaviour. The result of expecting precision to be high for such policies is the poor self-control seen in addicts. Low probability is assigned to competing action policies. This allows sensory information that predicts drug-use to exert an inflexible and dominating influence on their behaviour to the exclusion of other action policies (Pezzulo, Rigoli, & Friston, 2015; Schwartenbeck et al., 2015). In other words, too much confidence is placed in the policies that drive addictive behaviour, precluding an exploration of alternative affordances.

Now it might be thought that the PP theory of addiction we've just sketched is basically just a redescription of the reward learning theory. We will argue in the next section however that in the PP theory, dopamine should be understood in the context of the whole person in their engagement with the multiple relevant affordances of their econiche. Thus, addiction is not due to substances of addiction hijacking midbrain dopaminergic systems, even though we do not mean to deny the central role of such systems in the progression of addiction. Instead addiction

⁷ Technically, precision refers to the inverse dispersion of probabilistic beliefs. If the probability distribution is over a continuous variable, precision corresponds to the inverse variance. In predictive coding, this usually refers to the precision of prediction errors (Feldman & Friston, 2010). When the probability density is over discrete states, precision corresponds to inverse temperature; commonly encountered as a softmax parameter (Parr, Benrimoh, Vincent, & Friston, 2018; Parr & Friston, 2017). Precision correlates with the random fluctuations in a sensory signal. The more stochastic the signal the lower the precision. Our thanks to an anonymous reviewer for helpful suggestions on how formulate best the technical definition of precision.

⁸ This process of optimisation is given a neurobiological characterisation in terms of optimising the sensitivity of post-synaptic gain of cells. Phasic discharges in the dopamine system signal error in precision expectations. Tonic discharge of dopamine influences the post-synaptic gain on such error signals leading to an update of precision expectations (Friston, 2012: p.276).

should be understood in terms of how the whole person attunes to relevant affordances. What substance-induced dopamine does is *give the illusion of an agent attuning well to their environment when in fact in many cases the opposite is occurring*. The brain is best seen as an “organ of mediation” to borrow Thomas Fuchs's apt expression (Fuchs, 2017; Schütz, Ramírez-Vizcaya, & Froese, 2018). The brain mediates the relation of the agent as a whole to the environment, and it does so only as a part of a larger self-organising agent-environment system (Lewis, 2018).

4. An ecological-enactive account of habits

The PP theory we are developing suggests addictive behaviours are not passive reactions to sensory cues, as the reward learning theory would seem to imply. Action and perception are co-dependent and stand in a circular causal relation (Anderson, 2014; Clark, 2013). The agent acts with the aim of bringing about the future sensory states it expects to occupy. The sensory states the organism tends to bring about through its actions are those that it is likely to occupy over time if it is to remain in a state of dynamic equilibrium with the ecological niche it inhabits (c.f., allostasis). In substance addicts, the physiological states they come to occupy as a consequence of using the substance, are among those they learn to expect. The need for the drug can be compared to hunger. It is a physiological need that the agent must feed if it is to maintain a steady state, and remain in homeostatic balance with its niche. The urges and cravings that are felt in addiction are thus not external forces that act on the self from the outside (Schütz et al., 2018). They arise naturally as a part of the processes that sustain and nourish the agent the addict has become. Thus, the neurocognitive processes that contribute to addictive behaviour are not malfunctioning - they are not the product of a diseased brain. They are instead doing the work they should be doing for an agent that has become a substance-addict (Lewis, 2015).

Agents self-produce and self-maintain their identity as individuals over time, an organisational property of living systems referred to as “biological autonomy” (Di Paolo, Buhrmann, & Barandiaran, 2017; Maturana & Varela, 1980; Thompson, 2007).⁹ The “identity” of an agent as we will use the term goes beyond the continued existence or survival of the agent as a biological individual. It is the whole way of life of an individual agent that is produced and maintained over time by agents acting to fulfill their predictions. We can characterise this notion of identity in PP terms by equating the identity of an agent with the generative model an individual develops through its practical engagement with the world. In PP the agent is hypothesised to develop a hierarchically organised internal model of its bodily abilities in relation to its environment. This internal model is referred to as a “hierarchical generative model” because it is used to generate predictions of incoming sensory input over multiple spatial and temporal scales. Instead of building up an internal reconstruction of the world bottom-up, based on incoming sensory information, the brain is cast as pro-actively predicting incoming sensory input. It is this process of acting to bring about its own predicted sensory input that is referred to as “action oriented predictive processing” (Clark, 2013, 2015b), or “active inference” (Friston, FitzGerald, Rigoli, Schwartenbeck, & Pezzulo, 2017; Pezzulo, Rigoli, & Friston, 2018).

In our Ecological-Enactive interpretation of predictive processing we follow Friston and colleagues in conceiving of the generative model as being the whole organism in relation to the ecological niche (Friston, 2011). Biological systems are characterised by a set of attracting states that must be continually revisited over time if the system is to remain viable, and continue to exist. The set of attracting states can be characterised as an individual's way of life - they will be a function of its

⁹ For excellent discussions of this notion of biological autonomy in relation to PP see Allen & Friston, 2016; Kirchhoff, 2016; and Gallagher, 2017.

morphology, physiology, behavioural patterns and the ecologic niche it inhabits (Kiverstein, 2018; Ramstead, Badcock, & Friston, 2018). The model the agent comes to embody over time should ensure that it continuously revisits the attracting states that define it as an agent. In short, the model should sustain the way of life of the agent over time. Deviations from such an attracting set of states will be surprising to the agent, and will potentially threaten its existence. If the agent is to continue to exist over time, the model should steer the actions of the agents so that on average it samples sensory states that are among its attracting set, and are thus expected. Any agent that succeeds in minimising prediction error in the long-run will also thereby maximise the evidence for its own continued existence. Thus, in maximising the evidence for a model the agent is thereby self-producing and maintaining its own identity as an individual agent. Prediction errors are a measure of the disattunement of internal and external dynamics (Bruineberg et al., 2018; Bruineberg & Rietveld, 2014). Prediction errors arise in response to a dynamically changing environment either because something changes on the side of the agent in terms of its bodily needs, concerns and interests, or on the side of the environment. Thus, the notion of prediction error minimisation is dynamic – it needs to be constantly achieved anew in response to the agent's evolving circumstances in a volatile environment.

Habits typically form when an individual repeatedly and regularly engages in an activity. The repetition of the regular pattern of activity becomes a part of who the individual is, and therefore a part of an individual's identity. Habits can thus be thought of as abilities for engaging in activities that contribute to the production and sustaining of an individual's identity. To put this in the terms of PP, habits are we suggest abilities for avoiding unexpected sensory states. Anything that is unexpected is a threat to the continued way of life of the individual, and is therefore bad from the agent's perspective and should be avoided. Things that contribute to the sustaining of this way of life are good from the agent's perspective, and are therefore attractive and worth pursuing. An agent's capacity for regulating its coupling to the environment thus derive in part from its habits (Di Paolo et al., 2017). Habits we suggest are best thought of as abilities for maintaining adaptation to the agent's ecological niche, not as automatic behaviours set-off by sensory cues.

The habits the agent forms in addiction, and the rituals they routinely perform in using the drug can in extreme cases take over the agent's life. Their social life - the friends they meet, their work life, their relationship with partner and family - may gradually become organised around the sustaining of the way of life of the drug addict. In many cases the individual that the addict becomes is one that makes perfect sense given the challenges they face in their everyday life. Substance use has predictable, reliable effects in the life of a person that would otherwise face numerous physical, economic and psychological challenges. As Lewis (Lewis, 2018) notes, "exposure to physical, economic or psychological trauma greatly increases susceptibility to addiction" (p.1551). The habit of substance or alcohol abuse the individual develops can be thought of as growing naturally out of the life of a person otherwise fraught with difficulties (c.f. Pickard, 2012).

Our claim that addiction should be seen as a part of a person's identity is one we share in common with Owen Flanagan (see e.g. Flanagan, 2018). The habit of using a drug is, Flanagan argues, "identity conferring and identity constituting" (op cit, p.78). Addicts often engage in rituals of consuming a substance or drink as members of communities, and their individual identity is bound up with their sense of belonging to this community. Their addiction grows out of their participation in a way of life that confers on them their sense of who they are. Drinking alcohol for instance is "an extremely important feature in the production and reproduction of ethnic, national, class, gender, and local community identities...it is a key practice in the expression of identity" (Wilson, 2005, p.3, quoted by Flanagan, 2018, p.81). It is as members of a community that they are initiated into the behaviour of using a substance. Substance-use signals membership of a

community to which the individual values and wants to belong. Use as sanctioned by the community tips over into abuse once addiction takes hold of a person. The addict transgresses what the community regards as normal or acceptable. Through their actions of concealing, stealing, dissembling they gradually become a person they do not want to be. The reason it can be so hard for a person to break out of an addictive pattern of behaviour according to Flanagan is that the person's very identity is bound up with their way of life as an addict. To change requires them to literally become a different person.

We've been arguing that habits are identity-defining - they define who the person becomes over time, but they do not however settle the person's identity once and for all going into the future. Most people that go through a period of substance addiction succeed in escaping their addiction by their mid-30s, often without any professional help (Heyman, 2009; Lewis, 2015; Pickard, 2012). People are also motivated to abstain from using drugs if their careers require them to undergo random testing for substance use, as is the case for example with airline pilots (Heyman, 2009; Holton & Berridge, 2013; Lewis, 2015).

There is no denying however that for many individual's drug habits turn out to be *bad* habits in the long run, in the sense that they turn out to be a threat to the individual agent's identity. The addict no longer takes care of the social relationships and other projects that were formerly important to them. Their behaviour is also literally "self-destructive" because they no longer care for themselves. What at first sight seems like a healthy, adaptive response to a challenging life turns out in reality to be a retreat from life. The retreat from life can be the appropriate response for an agent embodying a model of the kinds of psychologically challenging environments addicts often tend to inhabit. However, it is the model of the environment the agent embodies, and the expectations that the agent forms on the basis of this model, that we will suggest can turn out to be harmful to them in the long-run.

5. Why is addiction harmful?

The way of life of the addict leads them to develop a generative model that is skewed towards dealing with the particular challenges of feeding their habit. The individual, through their activities, contributes to the construction of a niche organised around their drug habits. Their habits allow them to remain well-attuned, and keep prediction errors under tight control so long as they remain within the narrow confines of such a niche. However, the model that addicts come to embody is typically not well suited to the sorts of volatile environments we inhabit. Drug habits ultimately prove too narrow, and overly rigid and inflexible to maintain attunement to an ecological niche in flux. The addict for instance may lose their job, or find their business failing. Their grip on the niche in which they are situated proves to be too precarious to be sustainable in the long-run. Feeding their drug habit is destructive of who they were before becoming an addict, and an accidental overdose may of course even deprive them of their life.

Proponents of the incentive salience model have argued that addiction has some of the key characteristics of a mental illness because of the decoupling of desire from the agent's judgement of the best thing to do (Holton, 2009: ch. 5; c.f. Berridge, 2017) The agent might explicitly judge they ought to no longer use the drug, yet when encountering cues that predict the use of the drug they find themselves overwhelmed by temptation and their resolve is undermined. Levy (2014, 2019) has argued along similar lines, that what is dysfunctional, and aberrant in the behaviour of addicts is the way in which addicts can rapidly change their mind from judging at one time that they ought, all things considered, to stop using drugs, to judging that they ought perhaps to use them just this one more time (Op cit, p.338). Levy has provided an account of this "judgement shift" in addiction in terms of PP. The shift to judging that the drug is best consumed is the brain's way of explaining away the prediction error elicited by sensory cues that have come to be associated with drug use. The judgement that the drug should be used is thus a part of the model that does the best job of

explaining away current prediction errors. Levy takes the oscillation in the addict's judgements to arise as a response to prediction errors encoded by the dopamine system.

Levy takes dopamine to signal a prediction error, along similar lines to reward learning theories in which dopamine signals reward prediction error. We've suggested by contrast and in line with Friston and colleagues that dopamine is involved in weighing the precision (the reliability and salience) of relevant affordances (Friston, 2012; Friston et al., 2014, 2012; Schwartenbeck et al., 2015). When dopamine neurons fire in response to sensory cues predicting drug use, this is because those sensory cues are associated with high precision affordances. The sensory cues predicting drug use are not unexpected on our reading of PP - they do not elicit prediction errors. On the contrary, they are cues the agent selectively samples, and has purposefully sought out because they predict the agent is likely to succeed in the future in occupying the sensory states they expect. In order for sensory cues to elicit judgement shift they would have to give rise to prediction errors. But we suggest what Levy has missed is the role of action in fulfilling prediction. The sensory cues that are associated with drug use such as finding yourself in a neighbourhood where drugs can be scored, or among a group of friends that use a drug, are sensory cues the individual actively engenders because they are already well predicted, and the agent acts to fulfill its predictions. Dopamine signals the agent's high degree of confidence that its current sensory states will lead to expected future sensory states (i.e. those consequent upon drug use). Drug-use has become for the addict a self-fulfilling prophecy.

We argue that what proves maladaptive for the addict is the high precision assigned to affordances related to drug seeking and taking behaviours. This weighting of precision increases the probability of selecting these sorts of affordances as inviting action. On this view, one would therefore distinguish good from bad habits in terms of precision (i.e., the confidence placed in affordances) – not in terms of judgement shift as Levy proposes (Levy, 2014). In other words, bad habits are simply policies that are not fit for purpose in a volatile world; however, they are selected repeatedly because alternative options are not entertained. Although addictive policies may be good for drug-taking, they may be bad for everything else, especially if the social econiche (or therapeutic support system) does not support addictive behaviour. In a volatile environment, what the agent should do to remain well-attuned is to modulate the precision assigned to affordances in response to such volatility. An inability to do this will lead to suboptimal action selection, with potentially devastating consequences.¹⁰

In the next section we will show how the tuning of precision-weighting should be thought of as a process that takes place in the whole body, not only in the brain. The agent is able to remain well-attuned to a volatile environment by making use of feedback from the body about the rate of error reduction to set precision on relevant affordances. This addition, as we will see, will help further our view of addiction as a breakdown not of the brain but of the wider agent-environment system.

6. Weighing precision in the body

An agent that weighs the precision of relevant affordances would, we suggest, benefit from tracking the rate at which error is accumulating or decreasing over time. We use the term “error dynamics” to

¹⁰ Exactly the same mathematical processes can be used to describe selection for selectability. For example, fruit flies increase their mutation rate when exposed to volatile (temperature) environments. Conspicifics that are unable to adjust their mutation rates (i.e., precision) have a suboptimal encoding of environmental volatility and are outcompeted by conspecifics that can explore different phenotypic options. Our thanks to an anonymous reviewer for this comparison, and for comments on the difference between good and bad habits that helped us to refine our formulations in this section.

refer to changes in the overall rate at which errors are accumulating or reducing over time. “Error” here refers to divergence from future sensory states expected as a consequence of acting. The rate of change in error reduction thus refers to how fast or slow the sum of prediction error is being reduced over time relative to what was expected (Kiverstein, Miller, et al., 2019). The expectations in question are examples of what we've been calling “precision expectations”. They are expectations about the likelihood of a relevant affordance leading to expected future sensory states. If the speed of error reduction increases, this equates to a faster reduction in prediction error over time relative to what was expected. This feedback should then act as evidence for an expectation that a relevant affordance has high precision. If speed of error reduction decreases, this equates to an accumulation in prediction error over time. It indicates that a relevant affordance is failing to lead to expected future sensory states. The accumulation of prediction errors should therefore lead to a decrease in precision.

The performance of an action policy in reducing error can be plotted as a slope that depicts the speed at which errors are being accommodated over time. The steepness of the slope indicates that error is being reduced over a shorter period of time, and so faster than the agent expected. The steeper the slope, the faster the rate of reduction. If the speed of error reduction increases, this equates to a decrease in prediction error over time (relative to what was expected). In that case, the action policy should be weighed as more precise. If speed of error reduction decreases, this equates to an increase in prediction error over time, with the result that the agent fails to occupy the future sensory states it expects. Feedback of this kind should be taken as evidence for weighing an action policy as having low precision.

We have recently suggested that error dynamics are registered by the organism as embodied feelings (Kiverstein, Miller, et al., 2019; c.f. Van de Cruys, 2017; Joffily & Coricelli, 2013). Positively valenced bodily feelings indicate better than expected error reduction. Negative valenced bodily feelings provide feedback that a policy has reduced error at a worse than expected rate. Suppose you are a smoker and you find yourself sitting through a long and somewhat tedious talk at a conference. You could have gone for a smoke before the talk but you opted instead to wait till the next break. The speaker is running over into the break, and you begin to experience a strong craving for the cigarette you promised yourself. The craving you experience is, we suggest, your body telling you that there is relevant source of error that you were expecting to soon reduce. This situation is felt in the body of the agent as an unpleasant feeling of error on the rise, or tension. This negative feeling may lead the agent to explore the environment for other alternative possibilities to smoking that reduce tension for the duration of the lecture. Relevant possibilities that might now stand out soliciting them to act may be possible distractions such as doodling, looking up your email on your phone or continuing work on the slides for your own talk.

Positive and negatively valenced feelings provide feedback on the quality of the organism's engagement with the environment (c.f. Polani, 2009). These feelings are embodied as part of the valuation process that works as a sort of bodily barometer, keeping the organism informed about how it is fairing in its attempt at maintaining its adaptedness to its niche (Barrett, 2017). Agent's are normally sensitive to the rise and fall in error reduction and make use of this information about how well they are doing overall in reducing error to learn precise policies.

We have seen above how in our Ecological-Enactive interpretation of PP dopamine scores confidence in relevant affordances. We've suggested the dopamine system would do well to make use of changes in the rate of error reduction. Thus the process of assigning precision to relevant affordances will work best, in part, through keeping track of error dynamics. In the reward learning literature (discussed briefly in Section 2), dopamine discharge signals a reward prediction error that indicates something better (or worse) than expected has just

happened.¹¹ We've seen how by contrast in PP rewarding outcomes are what are expected. We therefore hypothesise that what dopamine is actually signaling is how well or badly an organism is doing at bringing about the future sensory states it expects. This hypothesis would allow for PP to take advantage of insights from the reward learning literature. Information about rate of change in error reduction is valuable feedback that can be used to fine-tune precision expectations. Making use of such information will ensure that the agent is able to adapt its actions to a volatile environment in dynamic flux. An agent that acts on precise policies, in other words, shouldn't just be interested in error reduction but in whether error has been reduced better or worse than expected.¹²

Furthermore, we have suggested that error dynamics are felt in the body in the form of positively or negatively valenced feelings. Thus, there is good reason to think precision expectations might be updated in part through using feedback from the body (Pezzulo et al., 2018). The important take home message here is that the setting of precision weighting isn't just a brainy event. Precision expectations are tuned to the agent's changing circumstances based on bodily feelings. We will show next how drugs of addiction create the feeling of reducing error at a better than expected rate. However, as we will see, this is often only an illusory experience of error reduction rather than tracking an actual increase in attunement. The reality is often a steadily increasing disattunement with affordances that should be of significance and stand out as salient to the person.

7. Addiction as tending towards a *sub-optimal* grip

We've argued in the previous section that it *feels* good to the agent when prediction error is reduced at a faster rate than expected, and it will feel bad for the agent when error builds up they are unable to reduce. This is felt in the body of the agent as feedback that error reduction has gone better or worse than expected. These feelings we suggested play a part in setting the agent's precision expectations ensuring that the agent typically attunes to the affordances that are important to them. Now consider what happens in substance addiction as the use of the substance exerts a tighter and tighter grip on the person. The drug acts either directly or indirectly on the dopamine system that is weighing the precision of relevant affordances. Crucially, it *provides feedback that error has been reduced at a faster than expected rate*. Each time the agent uses the drug, the same feedback occurs. Instead of their expectations simply being met, which would normally signal to the brain nothing new to be learned here, the brain responds by producing dopamine that signals that there is still something new and surprising to be learned. The reward learning theory claims that dopamine neurons signal unexpected reward: the agent has done better than expected in relation to the rewards that were expected. On the view we are developing, dopamine neurons signal the degree of confidence in predictions of sensory outcomes. We agree dopamine tells the organism that something better than *expected* just happened. However, the expectations that dopamine underwrites relate to the precision of relevant affordances - the predictions of the attainability of its future expected sensory states. The policy of using the drug has led to error reduction, and thus to the future sensory states the agent expects at a faster rate than was expected. Thus, the effects of the drug on the brain are to confirm the expectation that precision for the affordances of substance use should be set high.

¹¹ Dopaminergic neurons in the midbrain fire together at a rate that is correlated to the organism's expectation of reward. An increased or decreased firing rate indicates that an outcome that was better or worse than expected has occurred (Schultz et al., 1997).

¹² Additional evidence comes from the brain network that has been hypothesised to play a role in tracking error dynamics. Joffily and Coricelli (2013) suggest that orbitofrontal cortex in collaboration with the striatum are likely to be candidates for implementing processes that calculate rate of change in error reduction (p.13).

Each time the agent acts on a drug-using policy, they predict sensory cues that are associated with the pursuit of the policy, such as being in a particular neighbourhood where you can score the drugs you are seeking. These predictions give rise to prediction errors, which the agent acts to reduce by sampling the world in search of the predicted cues. The prediction errors arise from an affordance that is given high precision, and thus get to drive behaviour to actively seek out those cues. The agent will follow the paths of action (action policies) that are most likely to lead them to the drug. Thus, in addition, one is not automatically triggered to act by bottom-up sensory cues. The addict pro-actively seeks out the sensory states they expect, and in turn is directed to act in ways that will fulfill their expectations.

Repeated use of drugs of addiction can thus be thought of as training expectations for error reduction at a certain rate. Importantly, this is the source of positive feelings that comes with drug use (at least in the early days) on our account. Drugs of addiction act directly on the system that is signalling the probability that a policy does a good job of reducing prediction error. Thus it makes it seem there is now something they can do that is a failsafe, guaranteed means of arriving at the sensory states they expect and value. Addictive substances make it appear to the agent as if error is being reduced rapidly at a rate that is faster than anything the agent has anticipated. As soon as the drug wears off, prediction errors begins to increase again. Nothing was in fact resolved in the world through taking the drug though: there was only the illusion of error reduction. In fact, the addict often finds themselves in a worse situation, as is reflected in the negative affect associated with feelings of guilt and shame in the short term and loss of health in the long term.

Cravings in the addict can be thought of as the result of the accumulation of error - much like in the smoking example discussed above. The strong drives or cravings they experience are, we suggest, due to expectations of fast error reduction. As long as they don't use the drug, error accumulates that seems to be quickly resolved by finding and taking the substance. The addict feels bad so long as they are not using the drug because they are failing to meet the slope of error reduction they have come to expect through using the drug (cf. Koob and Le Moal, 2001, 2005). They are therefore driven to use the drug again in order to meet the rate of error reduction they have come to expect. Thus, the cycle of seeking and using takes hold and exerts a tighter and tighter grip on the agent.

Moreover, the addict has now come to expect a certain rate of error reduction - they have come to expect to do well at avoiding surprise even though the reality may confront them with many challenges and frustrations they are unable to manage. Consider a person who is constantly facing hunger because they don't have the money to buy food, and is cold because they are unable to pay to heat their homes. They expect to be well-fed and to stay warm, but their socio-economic status means that meeting these expectations is a continuous struggle. People faced with such difficulties in life struggle to meet their expected slope of error reduction and might be more attracted to the possibility to "self-medicate", as it is sometimes described. Once they have discovered the possibility to reduce disattunement with the world in ways that otherwise prove a struggle, one can imagine the temptation to do so repeatedly might be high. Marc Lewis makes this point well in relation to the susceptibility of people struggling with PTSD and depression to addiction:

"Importantly, it's not just attraction or desire that fuels feedback loops and promotes neural habits. Depression and anxiety also develop through feedback. The more we think sad or fearful thoughts, the more synapses get strung together to generate scenarios of loneliness or danger, and the more likely we are to practice strategies—often unconsciously—for dealing with those scenarios. Neural patterns forged by desire can complement and merge with those born of depression or anxiety. In fact, that's a lynchpin in the self-medication model of addiction. Gabor Maté persuasively shows how early emotional disturbances steer us toward an intense desire for

the relief provided by drugs (Maté, 2008), and Maia Szalavitz vividly portrays her experience as a late adolescent trying to brighten her depression with cocaine and ease her anxiety with heroin (Szalavitz, 2016). So, when we examine the correlation between addiction and depression or anxiety, we should recognize that addiction is often a partner or even an extension of a developmental pattern already set in motion, not simply a newcomer who happened to show up one day” (Lewis, 2017, p. 10).

What might seem like a fool proof way to reduce uncertainty is in fact no such thing. Addicts choose the familiar option, and continue to do so even when the outcomes are negative. They do not gather more evidence that might lead them to change their behaviour. The possibility to explore and find a different means of maintaining adaptation to the environment is down-weighted relative to the option of continuing to exploit the known consequences of using the substance. We can think of the behaviour of the addict in terms of a dynamical landscape of attractors. Typically, the attractor landscape changes over time as environmental conditions and the agent’s needs and interests change (Friston, Breakspear, & Deco, 2012; Rietveld, Denys, & Van Westen, 2018; Rietveld & Kiverstein, 2014).¹³ This is necessary if agents are to maintain a good grip on a volatile environment. They must sometimes not just stick to what they know, always doing what is already well-learned, but instead explore just-uncertain-enough environments that allow them to do better in the long-run in their dealings with a world in flux.¹⁴ The dynamical landscape of the addict is however made up of fixed-point attractors that do not destroy themselves over time. Instead they entrain the behaviour of the agent in ways we have just describing, locking them into rigid, and ultimately self-destructive cycles of behaviour (see also Friston, 2012; Schwartenbeck et al., 2015).

We suggest the reason addicts don’t explore and gather new evidence may be that substances of addiction make the person feel (at least temporarily) like they are well-attuned even though they are not. Such *drugs create an illusion of attunement to the environment*. Sensitivity to error dynamics is one way that good habits are woven into our skillful engagement incrementally over time - all directed to what matters to the organism. Drugs of addiction, as we have seen, lead the system to self-organize in relation to the environment in ways that lead agents to neglect the many other things in their lives that also matter to them in favour of the policy of feeding their addiction. The very same mechanisms that normally produce curiosity and exploration, when perturbed by addictive substances, produce precisely the opposite effect. Instead of being moved to pursue the multiple possibilities people typically care about, the addict find themselves increasingly being gripped by the drug infused field of affordances (Bruineberg & Rietveld, 2014; Rietveld, et al., 2018; Gibson, 1979).

8. Why addiction isn’t just a brain disease, and why it matters

The account of addiction we’ve proposed avoids falling onto one side or other of the dichotomy in which addiction is seen either as a biological disorder or as a purely social phenomenon whose causes lie for instance in poverty or in the urban environment. All accounts of addiction stress the importance of recognising the complex suite of

¹³ Clark has recently written, “Friston suggests, our ‘neural expectations’ may come to include expectations of ‘itinerant trajectories’ mandating change, exploration, and search. We ‘expect’ to sometimes engage in random environmental search as a means of entering into adaptively valuable states. To put it crudely, we randomly sample because - qua evolved organisms - we ‘expect’ to discover food, mates, or water at some point during the expedition” (2017: p.526).

¹⁴ Schwartenbeck, FitzGerald, Dolan, and Friston (2013) extend this direction of thinking by proposing that certain policies may be valuable insofar as they open the way the agent to visit multiple other states (Bruineberg & Rietveld, 2014; Rietveld & Kiverstein, 2014).

causes that lead up to addiction. There is now converging evidence for instance that physical abuse, economic inequality and injustice, and psychological trauma in early life increases the likelihood of addiction in the future (Satel & Lilienfeld, 2014; Sinha, 2008). The disease model of addiction acknowledge that social and environmental factors play an important role in the development of addictions impacting on the vulnerability and resilience of individuals. However, they often have an unfortunate tendency to downplay the agency of the addict, assigning too much importance to the brain. The contribution of the environment on such accounts is only to provide stimulation that passively drives the behaviour of addicts. The addict responds automatically to the stimulus properties of cues, their behaviour bypassing their conscious evaluation and control. The other side in the debate emphasises the social and historical causes of addiction, but in doing so downplays the importance of the brain in the development of addiction. Interestingly in common with the disease model these accounts also treats the agent as largely passive in the causal history that leads up to their addiction. The addict as an individual agent is passively acted on by their historical and social circumstances. Both sides in the debate fail to strike the right balance between explaining addiction in terms of its environmental and social causes, and explaining addiction in terms of its biological causes.

We suggest an account of addiction in terms of the dynamics of an agent-environment system self-organising in ways that minimise long-term prediction errors. On our account the dysfunctional behaviours of addicts are the result of *disorganization* within the agent-environment system as a whole. Human agents enter into a circular causal relationship with their surroundings. The organism’s perception of its environment, its actions and its feelings are co-determining. It is this dynamic relationship between the organisms and the environment that is disrupted in addiction. From this perspective addiction is best characterized not only as a change in particular neural circuitry, but as a more general loss of attunement of the organism and its environment. The resulting theory of addiction is thus one in which neural processes are necessary but not sufficient to account for addiction.

Explanations of addiction have been proposed by others in terms of PP that take the breakdowns associated with long term addiction to be a consequence of loss of contextualisation of how low-level habits by high-level processes of cognitive control (Clark, 2017, 2019; Friston, 2012; Pezzulo, Rigoli, & Friston, 2015). On those accounts the work of prediction error minimization maybe partly offloaded onto the body by allowing bodily habit to function as fast and simple heuristics that drive behaviour (Clark, 2015b). However, precision-estimation remains a function that is implemented entirely in the brain, and it is precision expectations that are held to account for the loss of contextualisation of habits by high-level cognitive control.

The view we have been developing by contrast takes precision weighting to arise in part out of processes that track error dynamics through bodily feelings in relation to the world. The learning of precision expectations on the basis of feedback from the body alters the dynamics within the organism as a whole so as to ensure that the organism remains well adapted to a dynamically changing niche. In other words, what changes when precision is weighted is not the encoding of precision expectations in the brain but how the organism and the environment fit together. Insofar as addictive substances impact core systems sensitive to error dynamics, these substances play a central role in altering how the organism and structuring of the environment continuously co-arise together.

Once we view addiction as a phenomenon of the whole agent-environment system, we can do justice to accounts of addiction that emphasise its societal causes (e.g Sullivan, 2018). We have argued it feels good to agents to be continuously improving in error reduction. Sometimes a person’s life however offers only the prospect of more uncertainty - think of soldiers that become addicted to substances while away in a strange land in a war situation. They can make a predictable and somewhat more comforting reality for themselves, out of what is otherwise the confusing reality of war, by means of substance abuse

because the substance can be trusted to have certain guaranteed and predictable physiological effects on the body. Once the soldiers return home to the predictable and familiar reality, drugs no longer present the attraction they once held. There are better policies available to the soldiers for improving attunement with the world. This may go some way towards explaining why rates of heroin addiction were high among soldiers stationed in Vietnam but upon returning home addiction rates fell back to their normal rates. The behavior of the soldiers stationed in Vietnam was in this respect somewhat similar to that of the rats in the famous *Rat Park* studies (Ahmed, Lenoir, & Guillem, 2013; Alexander, 2010; Alexander, Coombs, & Hadaway, 1978; Hari, 2015; Solinas, Chauvet, Thiriet, El Rawas, & Jaber, 2008). One group of rats were placed in simple cages all alone, but with plentiful opportunity to consume as much opioids as they wanted. For such a rat addiction was an inevitable outcome. When the same rats, now addicted to the substance, were moved to a much larger cage with other rats and a variety of games and opportunities for improving they tended to ignore the available opiates altogether (Alexander et al., 1978; Alexander, 2010). Given the current proposal, we think this could be explainable insofar as the rats were able to now meet their expected slope of error reduction, just like the soldiers returning home from Vietnam (Granfield & Cloud, 1999; Robins, 1993; Robins, Helzer, & Davis, 1975).

In addiction the agent is increasingly gripped by the environment until they cease to be open to the other non-drug related possibilities that may otherwise matter to them. Recovery from addiction then is likely to be facilitated by changing the expected rate (the slope) of error reduction itself through *restructuring (relearning) expectations for where to look in the landscape of affordances for error reduction*. A key part of undoing such habits we suggest will be developing new and different skills for reducing errors more efficiently, such as techniques of emotional regulation and mindfulness that help people to not only act in response to possibilities in the here and now, but also be open and give due consideration to engaging with possibilities that lie in the future (Garland, Froeliger, & Howard, 2014). One of the keys to escaping addiction may thus be *restoring openness to the many possibilities that matter to the agent*, and not only those that relate to their addiction.

9. Conclusion

In this paper we have argued that to understand what is harmful about addiction requires taking a wider vantage point on the organism-environment system as a whole. The brain of the addict is in fact doing what the brain is meant to do when viewed from the standpoint of the PP theory as we've interpreted it in this paper (c.f. Bruineberg et al., 2018; Kirchhoff & Kiverstein, 2019; Kiverstein, Miller, et al., 2019). It is continually optimizing the fit of the organism with its environment relative to what matters to the organism. Addictive substances make it seem to the organism as if error had been reduced but sadly for the addict this is just an illusion. The result in the long-run is almost inevitably a greater amount of uncertainty arising from a loss of sensitivity to the wider concerns of life.

If all predictive organisms care about is reducing error why isn't the life addicts lead at least one viable strategy for prediction error minimisation? Addicts become extremely skilled at organising their lives around the goals of finding and using the addictive substance. They develop models that are optimised to fit an environment in which these are the only things that matter. We've argued that typically predictive organisms don't only try and reduce error but reduce it at a particular rate. It might be thought however that this is exactly what the addict is doing as they get increasingly skilled at navigating an environment whose relevance is structured by their addiction.

What this misses however is the way in which all the drug can deliver is short-term reduction in error. The life of many (but not all) addicts becomes increasingly chaotic in other regards. As soon as the drug's effect wears off, what they return to is a world offering all of the uncertainty that never really went away. So long as the addict is high, it

seems to them as if they are succeeding at maintaining grip on what matters to them. Once the drug wears off, they find reality is very different. Substance addiction has been likened to a single room with many paths that all in the end lead the addict back into the same room again. The room of addiction is however fraught with difficulties and dangers. The progressive loss of contact with the rest of what matters can lead long term addicts to struggle with loss of material possessions and personal relationships, diminished self-worth, and physical health problems. Addiction can in this way lead to long term increases in error in relation to all the other things that matter to the addict. Humans have come to expect overtime to maintain relationships that matter to them, to hold onto their possessions, and to remain healthy. In addiction however they act in ways that frustrate these expectations. The point at which an addict in the end decides to really make a change is sometimes referred to as "rock bottom". This is the point at which what the addict has *actually* lost finally outweighs what they *feel* they are gaining.

The addict comes to embody a generative model that is tailored and built around the all-consuming activity of feeding their habit. What counts as an improvement of the model, its predictions and their fit with this environment is dictated not by finding a balance among the many things that matter to the addict. It is instead dictated to the agent by the increasingly wide range of possibilities that lead them back into the same vicious cycle of behaviour. What is harmful in the life of an addict is thus not to be found inside of the brains of addicts but in their wider engagement with life, and with the environment they enact.

Funding

Mark Miller carried out this work with the support of Horizon 2020 European Union ERC Advanced Grant XSPECT - DLV-692739.

Julian Kiverstein and Erik Rietveld are supported by the European Research Council in the form of ERC Starting Grant 679190 (EU Horizon 2020) for the project AFFORDS-HIGHER, the Netherlands Organisation for Scientific Research (NWO) in the form of a VIDI-grant awarded to Erik Rietveld, and by a project grant from the Amsterdam Brain and Cognition research group at the University of Amsterdam.

References

- Ahmed, S. H. (2004). Addiction as compulsive reward prediction. *Science*, 306(5703), 1901–1902.
- Ahmed, S. H., Lenoir, M., & Guillem, K. (2013). Neurobiology of addiction versus drug use driven by lack of choice. *Current Opinion in Neurobiology*, 23(4), 581–587.
- Allen, M., & Friston, K. J. (2016). From cognitivism to autopoiesis: Towards a computational framework for the embodied mind. *Synthese*, 1–24.
- Alexander, B. (2010). Addiction: The view from Rat Park. Retrieved July, 26, 2015.
- Alexander, B. K., Coombs, R. B., & Hadaway, P. F. (1978). The effect of housing and gender on morphine self-administration in rats. *Psychopharmacology*, 58(2), 175–179.
- Anderson, M. L. (2014). *After phrenology: Neural reuse and the interactive brain*. MIT Press.
- Arpaly, N., & Schroeder, T. (2013). *In praise of desire*. Oxford University Press.
- Barrett, L. F. (2017). *How emotions are made: The secret life of the brain*. Houghton Mifflin Harcourt.
- Berridge, K. C. (2017). Is addiction a brain disease? *Neuroethics*, 10(1), 29–33.
- Berridge, K. C. (2007). The debate over dopamine's role in reward: The case for incentive salience. *Psychopharmacology*, 191(3), 391–431.
- Bruineberg, J., Kiverstein, J., & Rietveld, E. (2018). The anticipating brain is not a scientist: The free-energy principle from an ecological-enactive perspective. *Synthese*, 195(6), 2417–2444.
- Bruineberg, J., & Rietveld, E. (2014). Self-organization, free energy minimization, and optimal grip on a field of affordances. *Frontiers in Human Neuroscience*, 8, 599.
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3), 181–204.
- Clark, A. (2015a). *Surfing uncertainty: Prediction, action, and the embodied mind*. Oxford University Press.
- Clark, A. (2015b). Radical predictive processing. *The Southern Journal of Philosophy*, 53(S1), 3–27.
- Clark, A. (2017). A nice surprise? Predictive processing and the active pursuit of novelty. *Phenomenology and the Cognitive Sciences*, 1–14.
- Clark, A. (2019). Beyond desire? Agency, choice, and the predictive mind. *Australasian Journal of Philosophy* (in press).
- Delgado, M. R., Miller, M. M., Inatt, S., & Phelps, E. A. (2005). An fMRI study of reward-related probability learning. *Neuroimage*, 24(3), 862–873.

- den Ouden, Daunizeau, J., Roiser, J., Friston, K. J., & Stephan, K. E. (2010). Striatal prediction error modulates cortical coupling. *Journal of Neuroscience*, *30*(9), 3210–3219.
- Di Paolo, E., Buhmann, T., & Barandiaran, X. (2017). *Sensorimotor life: An enactive proposal*. Oxford University Press.
- Elster, J. (1999). *Strong feelings: Emotion, addiction, and human behavior*. Cambridge, MA: MIT Press.
- Everitt, B. J., Belin, D., Economidou, D., Pelloux, Y., Dalley, J. W., & Robbins, T. W. (2008). Neural mechanisms underlying the vulnerability to develop compulsive drug-seeking habits and addiction. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *363*(1507), 3125–3135.
- Everitt, B. J., & Robbins, T. W. (2013). From the ventral to the dorsal striatum: Devolving views of their roles in drug addiction. *Neuroscience & Biobehavioral Reviews*, *37*(9), 1946–1954.
- Feldman, H., & Friston, K. (2010). Attention, uncertainty, and free-energy. *Frontiers in Human Neuroscience*, *4*, 215.
- FitzGerald, T., Dolan, R., & Friston, K. (2014). Model averaging, optimal inference, and habit formation. *Frontiers in Human Neuroscience*, *8*, 1–11.
- Flanagan, O. (2017). Addiction doesn't exist, but it is bad for you. *Neuroethics*, *10*(1), 91–98.
- Flanagan, O. (2018). Identity and addiction. *The Routledge handbook of philosophy and science of addiction* (pp. 97–109). Routledge.
- Friston, K. (2011). Embodied inference: Or I think therefore I am, if I am what I think. *The implications of embodiment (Cognition and Communication)*, 89–125 (Chicago).
- Friston, K. (2012). Policies and priors. *Computational neuroscience of drug addiction* (pp. 237–283). New York, NY: Springer.
- Friston, K., Breakspear, M., & Deco, G. (2012). Perception and self-organized instability. *Frontiers in Computational Neuroscience*, *6*, 44.
- Friston, K. J., Daunizeau, J., & Kiebel, S. J. (2009). Reinforcement learning or active inference? *PLoS one*, *4*(7), e6421.
- Friston, K. J., Shiner, T., FitzGerald, T., Galea, J. M., Adams, R., Brown, H., ... Bestmann, S. (2012). Dopamine, affordance and active inference. *PLoS computational biology*, *8*(1), e1002327.
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2017). Active inference: A process theory. *Neural computation*, *29*(1), 1–49.
- Friston, K., Schwartenbeck, P., FitzGerald, T., Moutoussis, M., Behrens, T., & Dolan, R. J. (2014). The anatomy of choice: Dopamine and decision-making. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *369*(1655), 20130481.
- Fuchs, T. (2017). *Ecology of the brain: The phenomenology and biology of the embodied mind*. Oxford University Press.
- Gallagher, S. (2017). *Enactivist interventions: Rethinking the mind*. Oxford University Press.
- Gardner, E. L., & David, J. (1999). The neurobiology of chemical addiction. In J. Elster, & O. Skog (Eds.). *Getting hooked: Rationality and addiction* (pp. 93–136). Cambridge: Cambridge University Press.
- Garland, E., Froeliger, B., & Howard, M. (2014). Mindfulness training targets neurocognitive mechanisms of addiction at the attention-appraisal-emotion interface. *Frontiers in Psychiatry*, *4*, 173.
- Gibson, J. J. (1979). The theory of affordances. *The ecological approach to visual perception* (pp. 127–143).
- Goldstein, R. Z., & Volkow, N. D. (2011). Dysfunction of the prefrontal cortex in addiction: Neuroimaging findings and clinical implications. *Nature Reviews Neuroscience*, *12*(11), 652 (Chicago).
- Granfield, R., & Cloud, W. (1999). *Coming clean: Overcoming addiction without treatment*. NYU Press.
- Gugliotta, G. (2003). *Revolutionary thinker* (August 21)The Washington Post.
- Hari, J. (2015). *Chasing the scream: The first and last days of the war on drugs*. Bloomsbury Publishing USA.
- Heyman, G. M. (2009). *Addiction: A disorder of choice*. Harvard University Press.
- Holton, R. (2009). *Willing, wanting, waiting*. Oxford University Press.
- Holton, R., & Berridge, K. (2013). Addiction between compulsion and choice. *Addiction and self-control: Perspectives from philosophy, psychology, and neuroscience* (pp. 239–268). New York: Oxford University Press.
- Joffily, M., & Coricelli, G. (2013). Emotional valence and the free-energy principle. *PLoS computational biology*, *9*(6), e1003094.
- Keramati, M., & Gutkin, B. (2013). Imbalanced decision hierarchy in addicts emerging from drug-hijacked dopamine spiraling circuit. *PLoS one*, *8*(4), e61489.
- Kirchhoff, M. D. (2016). Autopoiesis, free energy, and the life–mind continuity thesis. *Synthese*, 1–22.
- Kirchhoff, M. D., & Kiverstein, J. (2019). *Extended consciousness and predictive processing: A third wave view*. Routledge.
- Kiverstein, J. (2018). Free energy and the self: An ecological–enactive interpretation. *Topoi*, 1–16.
- Kiverstein, J., Miller, M., & Rietveld, E. (2019). The feeling of grip: Novelty, error dynamics, and the predictive brain. *Synthese*, *196*(7), 2847–2869.
- Kiverstein, J., Rietveld, E., Slagter, H. A., & Denys, D. (2019). Obsessive compulsive disorder: A pathology of self-confidence? *Trends in Cognitive Sciences*, *23*(5), 369–372.
- Koob, G. F., & Le Moal, M. (2001). Drug addiction, dysregulation of reward, and allostasis. *Neuropsychopharmacology*, *24*(2), 97.
- Koob, G. F., & Le Moal, M. (2005). Plasticity of reward neurocircuitry and the dark side of drug addiction. *Nature Neuroscience*, *8*(11), 1442.
- Levy, N. (2013). Addiction is not a brain disease (and it matters). *Frontiers in Psychiatry*, *4*, 24.
- Levy, N. (2014). Addiction as a disorder of belief. *Biology & Philosophy*, *29*(3), 337–355.
- Levy, N. (2019). *Addiction. The Routledge handbook of philosophy and science of addiction*. Routledge.
- Lewis, M. (2015). *The biology of desire: Why addiction is not a disease*. UK: Hachette.
- Lewis, M. (2017). Addiction and the brain: Development, not disease. *Neuroethics*, *10*(1), 7–18.
- Lewis, M. (2018). Brain change in addiction as learning, not disease. *New England Journal of Medicine*, *379*(16), 1551–1560.
- Linson, A., Clark, A., Ramamoorthy, S., & Friston, K. (2018). The active inference approach to ecological perception: General information dynamics for natural and artificial embodied cognition. *Frontiers in Robotics and AI*, *5*, 21.
- Maté, G. (2008). *The realm of hungry ghosts*. Toronto: Vintage Canada.
- Maturana, H. R., & Varela, F. J. (1980). Autopoiesis: The organization of the living. *Autopoiesis and cognition: The realization of the living*. vol. 42. *Autopoiesis and cognition: The realization of the living* (pp. 59–138).
- Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive hebbian learning. *Journal of Neuroscience*, *16*(5), 1936–1947.
- NIDA. (2018, July 2). Media Guide. Retrieved from <https://www.drugabuse.gov/publications/media-guide> (on 2019, November 26).
- Parr, T., Benrimoh, D. A., Vincent, P., & Friston, K. J. (2018). Precision and false perceptual inference. *Frontiers in Integrative Neuroscience*, *12*, 39.
- Parr, T., & Friston, K. J. (2017). Uncertainty, epistemics and active inference. *Journal of the Royal Society Interface*, *14*, 20170376.
- Pezzulo, G., Iodice, P., Barca, L., Chausse, P., Monceau, S., & Mermillod, M. (2018). Increased heart rate after exercise facilitates the processing of fearful but not disgusted faces. *Scientific Reports*, *8*(1), 398.
- Pezzulo, G., Rigoli, F., & Friston, K. (2015). Active Inference, homeostatic regulation and adaptive behavioural control. *Progress in Neurobiology*, *134*, 17–35.
- Pezzulo, G., Rigoli, F., & Friston, K. J. (2018). Hierarchical active inference: A theory of motivated control. *Trends in cognitive sciences*, *22*(4), 294–306.
- Pickard, H. (2012). The purpose in chronic addiction. *AJOB Neuroscience*, *3*(2), 40–49.
- Polani, D. (2009). Information: Currency of life? *HFSP Journal*, *3*(5), 307–316.
- Ramstead, M. J. D., Badcock, P. B., & Friston, K. J. (2018). Answering Schrödinger's question: A free-energy formulation. *Physics of Life Reviews*, *24*, 1–16.
- Redish, A. D. (2004). Addiction as a computational process gone awry. *Science*, *306*, 1944–1947.
- Rietveld, E., Denys, D., & Van Westen, M. (2018). Ecological-enactive cognition as engaging with a field of relevant affordances. *The Oxford handbook of 4E cognition* (pp. 41). Oxford University Press.
- Rietveld, E., & Kiverstein, J. (2014). A rich landscape of affordances. *Ecological Psychology*, *26*(4), 325–352. <https://doi.org/10.1080/10407413.2014.958035>.
- Robins, L. N. (1993). Vietnam veterans' rapid recovery from heroin addiction: A fluke or normal expectation? *Addiction*, *88*(8), 1041–1054.
- Robins, L. N., Helzer, J. E., & Davis, D. H. (1975). Narcotic use in Southeast Asia and afterward. *Archives of General Psychiatry*, *32*(8), 955–961.
- Satel, S., & Lilienfeld, S. O. (2014). Addiction and the brain-disease fallacy. *Frontiers in Psychiatry*, *4*, 141.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*(5306), 1593–1599.
- Schultz, W. (2016). Dopamine reward prediction error coding. *Dialogues in Clinical Neuroscience*, *18*(1), 23.
- Schütz, C. G., Ramírez-Vizcaya, S., & Froese, T. (2018). The clinical concept of opioid addiction since 1877: Still wanting after all these years. *Frontiers in Psychiatry*, *9*.
- Schwartenbeck, P., FitzGerald, T., Dolan, R., & Friston, K. (2013). Exploration, novelty, surprise, and free energy minimization. *Frontiers in Psychology*, *4*, 710.
- Schwartenbeck, P., FitzGerald, T. H., Mathys, C., Dolan, R., Wurst, F., Kronbichler, M., & Friston, K. (2015). Optimal inference with suboptimal models: Addiction and active Bayesian inference. *Medical Hypotheses*, *84*(2), 109–117.
- Sinha, R. (2008). Chronic stress, drug use, and vulnerability to addiction. *Annals of the New York Academy of Sciences*, *1141*(1), 105–130.
- Solinas, M., Chauvet, C., Thiriet, N., El Rawas, R., & Jaber, M. (2008). Reversal of cocaine addiction by environmental enrichment. *Proceedings of the National Academy of Sciences*, *105*(44), 17145–17150.
- Sullivan, A. (2018, Feb 20). The Poison We Pick. *New York Magazine*. Retrieved from <http://nymag.com/daily/intelligencer/2018/02/americas-opioid-epidemic.html>.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*, Vol. 1, No. 1. Cambridge: MIT press.
- Szalavitz, M. (2016). *Unbroken brain: A revolutionary new way of understanding addiction*. St. Martin's Press.
- Thompson, E. (2007). *Mind in life*. Harvard University Press.
- Van de Cruys, S. (2017). Affective value in the predictive mind. In T. Metzinger, & W. Wiese (Vol. Eds.), *Philosophy and predictive processing: Vol. 24* Frankfurt am Main: MIND Group. <https://doi.org/10.15502/9783958573253>.
- Volkow, N. D., & Fowler, J. S. (2000). Addiction, a disease of compulsion and drive: Involvement of the orbitofrontal cortex. *Cerebral Cortex*, *10*(3), 318–325.
- Volkow, N. D., Koob, G. F., & McLellan, A. T. (2016). Neurobiologic advances from the brain disease model of addiction. *New England Journal of Medicine*, *374*(4), 363–371.
- Wilson, T. M. (2005). *Drinking cultures: Alcohol and identity*. Berg Publishers.
- World Health Organization (2004). *Neuroscience of psychoactive substance use and dependence*. World Health Organization Retrieved from https://www.who.int/substance_abuse/publications/en/Neuroscience.pdf.